

Chapter 28

Netcentric Security

A hacker penetrates the Web site of the Central Intelligence Agency and defaces it. The Chaos Computer Club writes an Active-X based Trojan horse that is downloaded to a user's PC and then enacts a funds transfer from Quicken. A disgruntled systems administrator commits acts of sabotage that causes a company to lose fifteen million dollars of business after he leaves. Hidden form fields present in numerous corporate Web sites leave sensitive information open to major security risks. These are a few examples of potential dangers in the netcentric computing environment.

Preserving security of information as it travels across the Internet, or even within an intranet, is complex. The Internet is a public resource accessible worldwide, comprised of heterogeneous nodes that are managed locally with minimal systemwide policy. However, businesses today rely on the Internet for the transfer of increasingly sensitive information. The interaction between diverse components (e.g., databases, operating systems, firewalls, routers, and application servers) makes it difficult to ensure that fundamental security requirements are met throughout the system. Implementing effective security in the netcentric computing environments often means finding and dealing with the weakest link in a large system of complex and dynamic links. However, the challenges are not insurmountable. By designing security into a netcentric solution and implementing the appropriate application, infrastructure, and procedural controls, security can be appropriately aligned with business risk.

Exhibit 1 illustrates several potential areas of weakness in a basic netcentric application.

The following explains the exhibit, starting on the right side and moving to the left:

1. The first vulnerability is from someone authorized to use the corporate network. After the insider attack, the next potential breach is through the corporate firewall. If the firewall is circumvented, a primary line of defense is lost, and attackers can launch a series of attacks against the internal network.
2. Assuming the firewall is secure, a second and potentially more likely target is the application (i.e., Web) server outside of the firewall. A number of possibilities exist, including unauthorized access to a

Potential Vulnerability Points

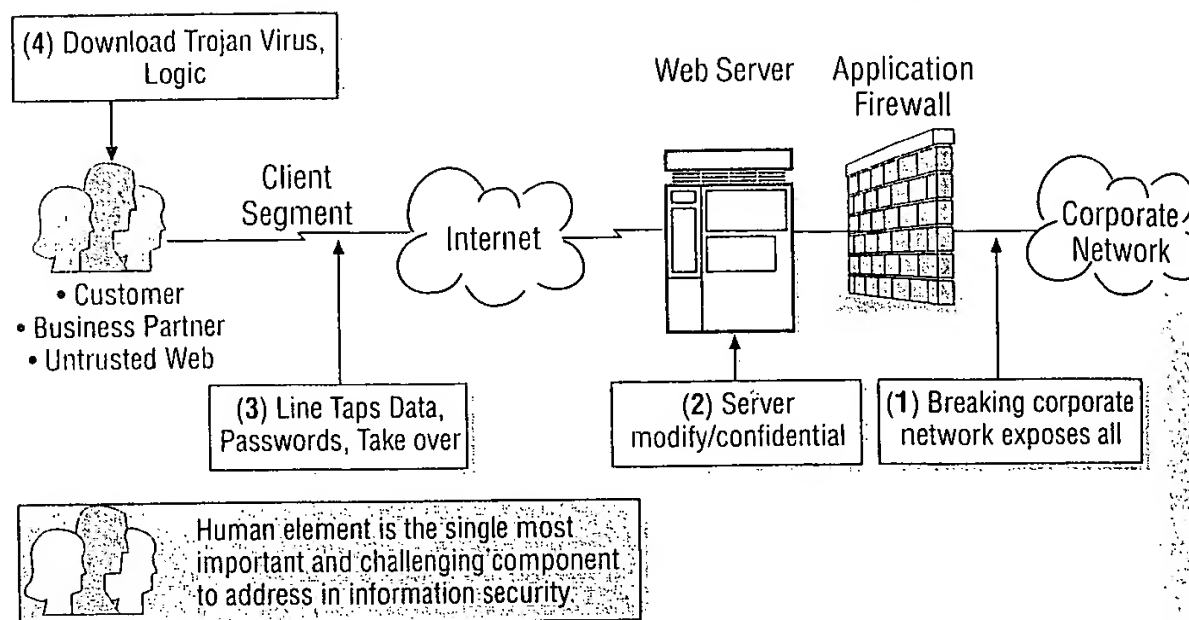


Exhibit 1. Vulnerabilities in Network Environments.

user account in order to withdraw funds from an investment account. A more serious breach may involve access to the operating system on the Web server: here, all user passwords can be intercepted, data can be captured and modified, and attacks may potentially be launched through the firewall against the internal corporate network and other supporting application servers, database servers, and other connected systems. Perhaps even more damaging is the use of such a Web server to represent the organization in a bad light or to conduct business transactions without the approval or knowledge of the organization's management. Liability for such transactions might be avoided, but the lasting damage to reputation may never be overcome.

3. A third risk involves interception of packets as they traverse the network between the client machine and the Web server. Fairly common are attacks on the Secure Sockets Layer (SSL) protocol, which is the protocol most often used to encrypt data across the Internet. Brute force attacks involving a number of computers in parallel have been successful when shorter key lengths, such as 40 bits, are used.
4. The most insidious technical risk of all is that of unauthorized code being downloaded to the user's computer and executing a harmful command. The Trojan horse written by the Chaos Computer club mentioned above is an example of this type of attack. Such code is acquired unwittingly either from sites run by individuals with malicious intent or from sites that are themselves victims of attack.




Exhibit 2. Layers of Security.

Although the sophisticated attacks, such as those against SSL, tend to get the most attention, it is actually the simple ones that are most often successful. The weak point is frequently a user who doesn't protect a password, a developer who doesn't code an application securely, or a system administrator who doesn't implement the necessary measures on a machine.

SECURITY LAYERS

The security objectives and requirements discussed in this chapter are applied to a variety of layers that exist in the physical computing environment. As shown in Exhibit 2, these layers are physical, network, operating system, database management system, application, and data.

The physical layer is the first line of defense. It represents basic controls like access to buildings and computer rooms. Physical security is one of the most critical layers since anyone with physical access to a machine may turn off any other security measures in place. Network security includes controlling access by firewalls, using subnets, and routing. Network security may be augmented by or include cryptographic assurances of correct data flow, data integrity, and confidentiality of data distribution. Operating system security has controls which provide authentication and access control services. The database and application layers provide additional controls on accessing data. One additional potential layer, not shown on this diagram, is middleware. Middleware products such as Tuxedo often

mediate access between the netcentric system and legacy applications. In the future, some of the greatest issues will be those associated with the relationships and interdependencies that are formed or that breakdown between these layers.

NETCENTRIC COMPUTING AND THE HEIGHTENED NEED FOR SECURITY

Networks are increasingly replacing individual computers, and access to all kinds of computing resources — such as CPU cycles, disk storage, and RAM — are being mediated by the network instead of individual boxes. This change is significant because the most pervasive computer networks and networking protocols were never designed to be secure. The following are several beliefs that IS managers have held:

- Security is done in the operating system (e.g., RACF and ACF2).
- Security primarily involves creating IDs and passwords, and is a relatively simple matter.
- We can implement the system with a low level of security and then increase it later.
- Data inside the corporate network are secure.
- We will start thinking about security when we begin to plan the system rollout.
- Security is managed by the “security group” and they often restrict access too much.

The following discussion looks at the important features of netcentric computing environments that lead to increased reliance on heightened security measures.

THE NETCENTRIC WORLD IS VIRTUAL

The Changing Meaning of Location

In the past, organizations relied on the physical security of a location as a means of controlling access to systems; this is still true to a lesser extent today. Access to the data center is strictly controlled. Physical access to PCs is controlled through building access security. Network security often involves the establishment of “trusted” network boundaries, which in the past have corresponded to a company’s officially designated places of doing business. However, telecommuting, virtual corporations, outsourcing, wireless networks, portable computers, and partnering arrangements are extending the boundaries of the network in increasingly complex ways. It is becoming increasingly difficult for security measures to rely on location as a means to validate the security access of a user.

Blurring Lines Between Internal and External Users

In the netcentric world, it is more difficult to distinguish clearly between internal and external users. Some of the greatest benefits in netcentric computing are achieved when we extend access to corporate systems to business partners and customers, creating more insiders. Why is this a security concern?

In testimony before the U.S. senate, FBI director Louis Freeh said, "A large portion of the computer intrusion reports that the FBI... receive have at their core an employee, or a former employee, who has exceeded his or her access, often in revenge for a perceived offense or wrong. These individuals have the knowledge of where the most sensitive information is stored, how to access the information, and, at times, how to steal or damage the data."

As more users become insiders, the levels of access and knowledge of systems expands towards a larger number of people with diminished loyalty to the organization. Consulting and outsourcing arrangements, telecommuting, vendor access, and virtual corporations all combine to make it difficult to draw concrete lines. Companies that maintain very strict security policies for external access will inhibit the formation of productive business partnerships and those who aren't careful enough may open up their network to intrusion from connected networks.

Less Permanence of Access

Permanence of access to systems is decreasing. Today, ad hoc partnering relationships are formed and then dissolved; an organization may have temporary needs for project resources, bringing in skilled people for short periods. Some consulting firms, for example, use development centers to support systems development activities. Several different development centers may be involved in a development effort, each communicating across a network with each other and with computing resources on the client's network. The project may be staffed at any combination of sites — either client or development center — and personnel may need varying levels of access to solution center and client systems for varying periods of time. After deployment, some of the maintenance activities may be outsourced and some of the company's employees may accept positions with the outsourcer. The project team access requirements must be met so that the job can get done, but access must be removed when it is no longer needed. In this transient environment, security must be flexible, timely, and address the complex issues that are present.

The Need to Establish Trust in Absence of Physical Observation and Contact

Most of today's retail purchasing involves a trip to the store and a face-to-face purchase. Customers know that the store is a legitimate store because

they can see it, walk into it, and carry out the physical merchandise. The store clerk will accept cash, validate a credit or debit card, or accept a check with validation of some form of identification. Even in the case of telephone-based catalog sales, customers are generally assured that they are talking to a bona fide merchant, because they initiated the calls to the well-publicized toll-free number, and the merchant is assured of payment by the credit card company.

In an eCommerce environment, it is more difficult to establish trust. The Web site that looks legitimate may exist only for the purpose of acquiring credit card numbers. Merchants and banks are affected because there is a lower level of assurance that the credit card number being provided is actually valid without the presence of a physical card. People are more likely to initiate fraudulent transactions when they do not come into physical contact with those whom they defraud.

Companies are rapidly discovering new types of consumer fraud that are unique to the netcentric environment. These types of fraud are not addressed by standard credit card authorization mechanisms or methods, such as Address Verification Screening (AVS), that have been used to combat fraud in the mail order environment. The stakes can be even larger in business-to-business transactions because the risk goes up as the number and value of those transactions increase.

ORGANIZATIONS ARE INCREASINGLY DEPENDENT ON KNOWLEDGE AND THE SYSTEMS THAT MANAGE IT

Increasing Importance of Knowledge Capital and Information as an Asset

The terms "information age" and information superhighway" are already cliches, even though we have only begun to realize the power of knowledge management in the organization. Information and knowledge will be primary currencies of this coming century, and organizations will need to manage information efficiently to compete effectively. One has only to observe the success of eCommerce enterprises to be convinced of the diminished importance of hard assets, such as plants and equipment, for some industries.

Knowledge management systems, data warehouses, and other types of information access will continue to increase in importance. A greater need to make knowledge capital and corporate databases available inside and outside the enterprise and the increasing value of that information serve to make protection of knowledge capital and information a significant challenge as organizations attempt to achieve a balance between security and the need to share information.

Increasing Dependence on Computers and Networks

Netcentric computing and client/server technology continue the trend toward increasing dependence on computer networks, as mission-critical applications are developed and new potential points of failure are created. When businesses use this technology to collaborate with one another, forming intricate relationships, this complexity is multiplied. When communication with vendors, partners, and customers depends on computer technology and when that communication becomes integral to the delivery of customer products and services, integrity and availability become the essential ingredient for success.

THE TECHNOLOGY IS EVOLVING WITH EVER-INCREASING SPEED

Increasing Intelligence of Devices and Communication with Those Devices

In the past, security was often under the purview of a small group of specialists who were experts in securing the centralized mainframe environment. With the move to client/server, this has been changing dramatically as more and more critical processing becomes distributed. Netcentric computing continues this trend with more complex functions and the corresponding distribution of security mechanisms throughout an organization, together with its business partners.

Furthermore, fewer than three percent of the microprocessors produced today go into traditional computers. The rest go into a host of other products and devices including medical equipment, appliances, stereos, environmental control systems, process controllers, navigation equipment, and network communications devices. As these products and devices become more sophisticated, more interactive, and ubiquitous, the paths of access among these devices and computers will be increasingly complex.

For example, Sun Microsystems unveiled a product called Jini, which uses the Java programming language to harness the power of potentially millions of computers, ranging from giant mainframes to tiny palm-sized devices. These kinds of developments, as one press report wrote, have "encouraged developers to proceed on the assumption that every home, car and other personal environment will eventually be part of — and empowered by — a universal network ... a world in which millions of small programs, called objects, seamlessly flit back and forth between tens of thousands of devices that have been enabled to recognize and be recognized by the network."

In a world of ubiquitous computing, security provisions must keep pace and yet not be too intrusive. Traditional ways of securing access, such as an ID and password at the device level, will not scale well in this environment.

Increasing Concerns Over System/User Privacy

The dark side of the visionary accounts of the world of ubiquitous computing is the growing concern for personal privacy. The ubiquitous network means, for example, that as we drive our cars, devices in the cars will communicate with systems to tell us our location, give us traffic and weather reports, and advise us on best routes. That capability also brings with it, however, the capability to be constantly tracked by "Big Brother." The speed with which data can now be aggregated and the use of focused intelligent agents to isolate data of interest means that everything from buying patterns to sensitive medical histories are possible targets.

Privacy is often confused with security; however, privacy is a somewhat different issue which is focused more on protecting information associated with individuals. Many countries have laws involving privacy of certain types of personal data such as medical records, addresses, and phone numbers. These laws and the rights of individuals will have increasing importance in the future.

Several concepts are important when considering privacy:

1. People should know what information is on file about them and should know how that information is used.
2. People should have the right to challenge and correct any erroneous information.
3. There should be no secondary use of the information without the person's consent.
4. The custodian of the information has an obligation to maintain security and quality of that data.
5. Some type of public notice should occur in cases where systems store private information about individuals.

Increasing Rate of Change in Technology and Its Deployment

Every new program and every new major enhancement to an existing program introduces the potential for software bugs. Some of the bugs may ultimately affect security. Examples of this have been found in browser software as well as in Windows NT and UNIX operating systems. In some cases, new technology is developed without sufficient consideration of the required security. In still other cases, the security features are there but people responsible for implementing them do not know how to use the security features. In other cases, people may know how to use the features, but they have difficulty in managing them across the plethora of environments.

It is common for IS managers to wait to expend resources on security because they are waiting for the technology to mature. However, in fact, products never really "mature"; they will just keep enhancing features and

functionality to keep pace with new challenges. In fact, it is critical to select solutions that will keep pace with new versions of the technology, new standards that are being developed, and new industry directions so those solutions will integrate well with products under development.

Increasingly Sophisticated Attacks on Security

Easy-to-use software tools that automate sophisticated attacks on networks circulate freely on the Internet, so these tools are accessible to a broad range of unskilled intruders. Some examples of these tools include sophisticated programs that take advantage of weak points of the TCP/IP protocol and perform activities such as taking over an established user's sessions (also called a session hijacking). Such attacks have always been possible, but the technical complexity of performing the attack has prevented widespread abuse. Commercial and public domain software such as Satan and password crackers can also be used to look for security holes on devices attached to networks. Although these tools serve a valid purpose in the hands of a responsible party, they can also be used by attackers.

One of the most significant issues with netcentric security involves weaknesses at the client (a PC, PDA, etc.). Client operating systems often do not implement strong security. The fact that codes and files are downloaded to the client from a source that is often untrusted, raises all types of security concerns. Although there have been some attempts to combat these issues in browsers and in popular programming environments such as Java, these methods are far from foolproof. Malicious applets can redirect the PC to a phony Web site, steal passwords and certificates, and even launch attacks against the internal network to which the client is connected.

In the future, some of the most serious attacks will be directed against weaknesses in the client, because it has some of the most serious inherent security weaknesses. In addition, the greatest vulnerability will lie in network management systems, system management platforms, and key management/key recovery infrastructures. These components become the "keys to the kingdom"; once compromised, they grant access to a great number of machines on the network.

WHY SOLVING THE SECURITY PROBLEM IS IMPORTANT

As this book stresses, netcentric computing leads to immense business opportunities, which include

- *Lower cost.* The cost of performing a transaction through electronic media is an order of magnitude lower than the cost for a face-to-face transaction.
- *Increased efficiency.* Businesses and consumers can be in touch with resources around the world at a moment's notice.

- *Speed to market.* Netcentric technologies are more compartmentalized and reusable than older client/server counterparts, and also facilitate knowledge sharing from developers around the world, radically reducing development time.
- *Improved customer service.* Netcentric technology goes beyond simply improving customer service. Through media like the Internet and kiosks, technology solutions enable customers to serve themselves
- *Creates new services.* Netcentric computing does more than permit new channels to offer services; it adds a service dimension to physical products. Netcentric computing enables new products that meet new market needs.

If security issues cannot be solved, the technology cannot be implemented safely. To realize the potential benefits of netcentric computing, security issues must be addressed in a thorough, dynamic, and flexible fashion. New threats and risks evolve quickly in the netcentric environment, and security programs will become ineffective and obsolete if not reviewed and updated regularly.

A FRAMEWORK FOR IMPLEMENTING SECURITY

Today's netcentric computing infrastructure requires a complex mix of operating systems, Web servers, database servers, firewalls, management tools, routers, and underlying network components. Each different component of this infrastructure has specific security considerations that need to be addressed. Each component supplies some level of security protection, at the same time offering a target to an attacker.

How can organizations bring a set of components together to form a system and then determine the security attributes of the system from the analysis of the properties of the system's components? Companies are faced, on the one hand, with a set of high-level, abstract, system requirements and, on the other hand, with a combination of commercial products that may be configured in a multitude of ways and special-purpose devices and/or software developed by the system integrator. There is a substantial intellectual gap between the high-level statement of requirements and the realities of the implementation.

To bridge this gap, companies must start with a set of *security policies*. A security policy is the set of rules, directives, and practices that regulate how an organization manages, protects, and distributes sensitive information. A security policy is translated into access control rules that are enforced by the system. The desired attributes of the environment are realized, in turn, by the implementation of a set of "mechanisms" — functions that can be shown to provide the requisite attributes. The critical point is that one proceeds from policy (i.e., a high-level statement of the desired

global properties or characteristics of the system) to a set of specific mechanisms.

Understanding fundamental security objectives and requirements is a prerequisite to the implementation of a good solution.

Policy Objectives

Many high-level policies begin by articulating broad objectives for the system under consideration. Frequently, these can be gathered under one of three general categories, sometimes referred to as "CIA":

- Confidentiality: data should be accessible by only those properly authorized.
- Integrity: systems, and the data stored on them, should be immune to unauthorized modification.
- Availability: systems should be immune to denial-of-service attacks and should be able to meet the service levels they were designed for.

As information systems have matured, the body of laws and corporate regulation governing their use has expanded. As a result, many policy statements include objectives concerning *laws and ethics*, which insure that network, system, and security operations function within applicable laws, regulations, mandates, licenses, contracts, and "codes of conduct."

Although technology-based mechanisms play an important role in the enforcement of security policies, day-to-day procedures and management vigilance by management are required to achieve these objectives. For example, operational procedures and disaster recovery plans help ensure availability. Manual controls may help to ensure the integrity of data at certain key points in processing.

Requirements

Policy objectives are met by aligning people, processes and technologies to meet four fundamental security requirements: Identification and Authentication, Access Control, Audit, and System Integrity, defined below.

Identification and Authentication (I&A). An identifier is a piece of data used to uniquely identify an entity in a transaction. Real-world examples of identifiers include a drivers license or a national identification number. Identifiers must possess the following characteristics:

- *Uniqueness.* Each entity must have one unique identifier. No two entities have the same identifier.
- *Association.* There must be some way to bind the identifier to the entity (e.g., tying a social security number back to an actual person.)




Exhibit 3. Synergy Between Process and Technology.

Exhibit 4. Authentication Methods

Authentication method	Examples
What the user knows	A secret password, PIN number, credit card number and expiration date, mother's maiden name
What the user has	An ATM card, credit card, smart card, private key stored on an encrypted file on a PC
What the user is	Biometric verification such as voice prints, iris scan, signature verification, thumb scan

I&A provides a good example of the necessary synergy between process and technology (see Exhibit 3). Policy may mandate, for example, that requests for new accounts may be issued only by Human Resources managers and that each request must be properly logged. Identifiers are issued to entities during part of a registration process that validates an entity's request to participate in a system, generates a unique identifier, binds that identifier to the requesting entity, and distributes the identifier to the now participant entity. The technology is invoked at several points to support this process.

Similarly, once participating entities have been registered, an authentication mechanism validates the identifier during a transaction. Authentication is the process that validates that the entity requesting access, whether that is a human or automated process, is the true owner of that identity.

Authentication is performed by three primary methods, by validating

- What the user/entity knows
- What they have, or
- What they are

Exhibit 4 describes examples of each of these methods.

Access Control. Once identity has been established, access control rules determine what resources the entity may use. In one frequently used model of secure computing, the entities of interest in a system are "subjects" and "objects." A subject is an active entity, loosely described as a program in execution, and the surrogate of a person. A subject has an identity and attributes. An object is a passive entity, usually a repository of information.

The goal of the access control requirement is to reliably mediate the access of subjects to objects. On each attempted access of an object by a subject, the system determines whether or not the access is to be granted. It does this by applying a set of access control rules along with information it has about the subjects and the objects.

Access Control is used to permit or deny a specific type of use of system resources. For example, a user may be authorized to access a resource, but only for reading. Access control can be used to arbitrate access to files, processes, operating system ports, application functions, database tables, portions of a network (such as through virtual or dedicated circuits and firewalls), and other types of resources. This is accomplished most frequently through the use of Access Control Lists (ACLs). An ACL for a resource specifies the user or group and the type of access permitted (read, write, etc.). ACLs may optionally include date and time restrictions and program restrictions.

A refinement of traditional access control is referred to as "Role based access control" (RBAC). RBAC associates a job function/role to a set of resources, and then assigns the user to a particular role. Therefore, for example, the role of junior bookkeeper may have read and write access to the petty cash account, but read-only access to the general ledger. The advantage of RBAC is that it facilitates the management of access control and prevents users from retaining access to data that is no longer needed as they move from role to role.

Resource access control may be either restrictive or permissive. Restrictive resource access control is based on the policy that "whatever is not explicitly authorized is denied." Permissive resource access control is based on the policy that "whatever is not explicitly prohibited is allowed." Each of these methods has a use. For network and firewalls, restrictive access control is commonly used. For most servers, permissive access control is the norm.

Audit. Auditing is used to record accesses to resources and may be implemented at a number of layers, including operating system, database, application, and middleware as well as in network devices such as firewalls and routers. Auditing is typically implemented in combination of these layers to allow reconstruction of events after a security problem is detected. Good logs should be searchable for known or suspected patterns of abuse, and should be protected from alteration. Logs can monitor a variety of data, including access times, user IDs, locations, actions the user performed, and whether or not those actions were successfully completed.

A widespread perception is that logging is a post hoc security measure; it only has benefits after disasters, or hackers, have struck, and even then




Exhibit 5. Security Framework.

it will only help assess the damage done. Logging *will* do those things, but good logging capability, wisely used, is like the TV cameras hiding behind the smoked plastic domes on the ceiling at the mall: unobtrusive but potent deterrents to roguish behavior. When people know that somebody *might* be watching, they tend to behave better. If the unwanted behavior is not the result of malicious intent, proper auditing allows administrators to determine the cause of the problem. It is an excellent “lessons learned” tool. Well-designed log strategies allow the employment of power forensic analytical tools for determining the source cause for poorly-behaving systems. As such, auditing is a strong risk mitigation mechanism.

Integrity. Integrity refers to the property that any system must have if it is to protect itself and enforce the security policy. Integrity is the assurance that a system’s implementation (or a component’s implementation) conforms to its design. Of the four requirements, it is the most nebulous but perhaps the most important. Systems breached by buffer overflows, faulty parameters, or attacks on improperly configured network ports have failed to meet the integrity requirement. Viruses constitute what is probably the best known attack on integrity. Such faults appear at the boundaries of a system, and must be removed by a thorough analysis of those interfaces, as well as the use of mechanisms use to ensure that system files are safeguarded.

SECURITY FUNCTIONS

The appropriate combination of people, processes and technology should be applied to implement security requirements across all layers of the physical environment to achieve the objectives described above. This is illustrated in Exhibit 5.

People: Security Roles and Responsibilities

Neither the most advanced technology, nor the most careful procedures, will serve to secure a netcentric environment if the people who use and manage it lack the understanding and training to do their job properly. The failure of people to do what is required to secure the system is the most common failure in security. Some of the most significant barriers to security are

- Mismatched customs, cultures, and values
- Conflicts of divisional objectives and strategies
- Poorly defined and implemented roles and responsibilities

All of these can derail the best security architectures, models, and plans.

Several actions should be undertaken when planning and executing netcentric efforts:

1. A program of training and security awareness should be undertaken to sensitize responsible parties to security issues and to their basic responsibilities. Because security is implemented in many different places, and a breach in only one can result in failures in security, a well thought-out communication plan is critical.
2. Roles and responsibilities for security should be identified for all facets of the design, build, and run phases. Mechanisms, responsibilities and budget for proper operation should be clearly delineated for each phase of the effort. In addition, responsibilities should be established at the infrastructure level to complete necessary processes (described in the next section).
3. Required organizational changes must be planned and executed to support the netcentric environment.

Process: Administration and Management

Several processes are essential to secure a computer environment (Exhibit 6). These include

- Policy development
- Risk assessment
- Plan and Build
- Administration
- Compliance
- Change management

Security Policy Development. It is imperative that a well-constructed statement articulate the organization's goals and policies with respect to the use of computers and networks and the protection of the information they contain.

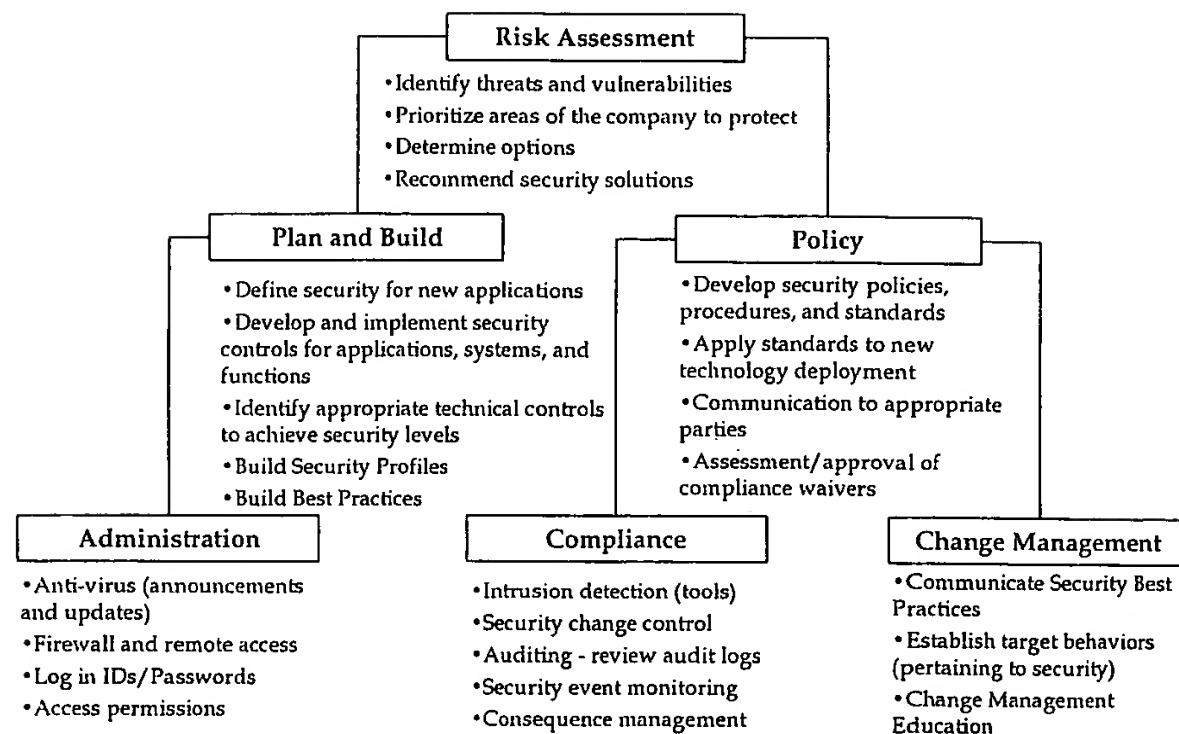


Exhibit 6. Essential Security Processes.

A number of issues must be addressed when developing a security policy. Example issues that may be covered include

- Who is allowed to use the resources?
- What is the proper use of the resources?
- Who is authorized to grant access and approve use?
- Who may have system administration privileges?
- What are the users rights and responsibilities?
- How is sensitive information classified?
- What happens when the policy is violated?

Supporting standards, guidelines, and procedures should be developed to provide direction to security implementation. These supporting documents may cover specific technologies and platforms. Change management helps to facilitate the adoption of the directions, and compliance processes are used to verify that the directions are being followed.

Risk Assessment. This is the process of identifying business risks, identifying system vulnerabilities or weaknesses that can affect those risks, and recommending mechanisms to control those vulnerabilities. The risk assessment process determines what is to be protected, something that is the cornerstone of any effort to secure a system. The assessor examines the organization's security policies and the infrastructure of processes and

mechanisms to find gaps in the fundamental security requirements. (If there are no policies in place, that is itself a risk.)

By determining what corporate assets are most in need of protection, the assessor gains a good understanding of the assets critical to the business. These assets can be

- Physical (e.g., systems, networks, facilities, customer data, or revenue stream)
- People (key employees)
- Intangible (e.g., reputation of the business)

The risk assessment involves a thorough appraisal of the environment from a technical, procedural, and business standpoint. The result is a set of measures designed to mitigate the risks in a cost-effective fashion.

Because no system can be made perfectly secure, the challenge of the risk assessment is to balance the tradeoffs between increased security and cost. Some assets may warrant large expenditures; others may not. A result of the assessment is a "residual" risk (meaning the risks that will exist after all recommended countermeasures have been implemented) that is consistent with the organization's needs and policies.

Plan and Build. This involves the integration of security into new technology design, development, and deployment processes. As described above, netcentric environments may provide control mechanisms in many different places, because security is integral to the application. For these reasons, it is essential to take the appropriate steps to design in security up front. In fact, the impact of implementing and maintaining security mechanisms can be so significant as to effect the entire value proposition of a system capability being considered.

If, for example, the solution will require costly smart cards and an infrastructure of readers, it may dramatically affect the cost per transaction. In addition to hardware cost, it is not uncommon to have significant resources on a project dedicated to implementing firewalls, developing public key infrastructures, integrating security into the application, locking down operating systems and Web servers, and implementing intrusion detection and audit tools. Such tradeoffs need to be considered carefully during the risk assessment.

It is also important to note that the Plan and Build environment is subject to business and technical risk itself. Consequently, it requires the development of an appropriate security policy to govern itself. If, for example, the test system is broken into and back doors are introduced, it may lead to later compromise of the production system.

Administration. This includes the processes to administer the environment securely, such as maintaining user accounts and security profiles, certificate issuance and revocation, configuring servers for proper security, and changing access rules. It is helpful if security administration is separately monitored. All administrative functions should be defined with proper attention to separation of duties to provide proper accounting and control of business processes. Metrics for service level monitoring should be collected and distributed to the proper accounting or monitoring function. Data collected for service level monitoring should be protected from tampering to ensure proper accountability and accuracy. Security administration must be separated from the Operations and Help Desk functions to provide segregation of duties and to ensure that the metrics of the organization responsible for the function are aligned with the need for security.

Technology: Security Mechanisms

Although rapidly changing technology represents a security challenge to the netcentric world, these same advances have provided valuable tools for the security architect. Among the most important of these new technologies is public key cryptography. Its chief merit lies in its ability to permit total strangers to communicate spontaneously and in secret without going through an elaborate preparatory process to establish secret keys. Without this form of cryptography, many of the benefits of netcentric systems could not be realized. The following sections provide an introduction to this important technology as well as the enabling technologies of Firewall, Audit Tools, and Intrusion detection software.

Cryptographic and Certification Services

Public key cryptography is one of the most important enabling technologies in the netcentric environment. Along with the Certification Services provided by a Public Key Infrastructure (PKI) this technology provides fundamental capabilities that are essential to the netcentric world:

- *Confidentiality.* As defined above, cryptography ensures that messages are accessible only by those properly authorized — even when they traverse insecure networks. (Note that the term “message” here can refer to an e-mail dispatch, or the more dynamic transactions of Web sessions.)
- *Authenticity.* This is the assurance that a message was actually sent by the purported sender.
- *Integrity.* This is the assurance that the message has not been modified in transit.
- *Nonrepudiation.* This is the assurance that a sender cannot disavow a message (see later section).

Cryptography relies on the use of "keys" to encrypt communications. There are two types of keys:

1. A *secret* key is shared between the two entities in a transaction. Because the same key is used to encrypt and decrypt data, this is referred to as *symmetric* key encryption. For the parties to communicate, they must establish the secret key in advance, using a secure channel. The most common implementation of a symmetric key algorithm is the Data Encryption Standard (DES.).
2. A *public/private* key pair or *asymmetric* key uses a pair of keys to encrypt and decrypt messages. Messages encrypted using one of the keys can only be decrypted with the other key. Each party possesses a pair of keys, one public key accessible to all participants in the system, and one private key accessible only to the party that owns it. The most common implementations of public key algorithms are supplied by RSA Data Security, Inc. In the most basic implementations, data are encrypted by the sender with the public key of the recipient and decrypted by the recipient with their private key.

Symmetric key systems can provide secure communications between two entities, but they have significant key management problems. "Key management" refers to those processes necessary to issue, maintain, and revoke keys as appropriate. Because a different symmetric key must be exist for each pairs of users, in a community of N participants there are approximately N^2 keys to track. Additionally, the life span of a symmetric key is short (1 day is common for a 56 bit DES key) so symmetric keys must be changed frequently. Tracking up to N^2 keys per day adds significantly to the complexity of key management.

Although public key cryptosystems do not require users to share a common secret key, key management is still a serious problem. Public key systems require a binding between a specific public/private key pair and an entity that is participating in the system. When using a public key to protect information destined for a specific entity, the user assumes that the public key he or she uses is really the one belonging to the entity. The only way to assure this binding is through the use of a trusted third party (TTP), called a "Certificate Authority," or CA.

Recall that the method for transmitting a message using public key cryptography is to encrypt the message with the receiver's *public* key. The benefit is that a user's public keys can be sent as clear text, or even published in a directory. Therefore, if Alice wants to send a message to Bob, but is tricked into using Eve's public key, then Eve will be able to intercept the message. (Eve can then, if she chooses, reencrypt the message using Bob's actual public key, and neither Alice nor Bob will be the wiser.) In a global

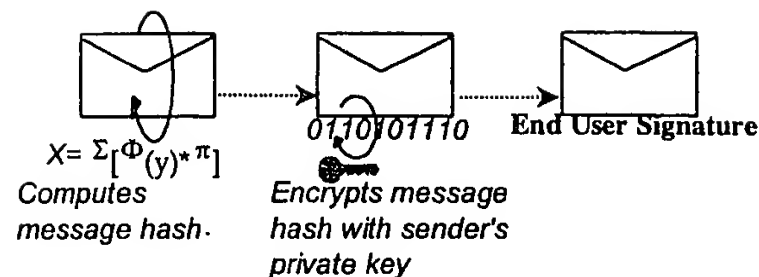


Exhibit 7. Digital Signatures.

network lacking face-to-face contact, users must be assured they are using the right key. The CA provides this.

The CA serves a function analogous to that of a passport or drivers license. The CA binds public keys to users and services similar to the way a passport agency issues you a passport that ties your name and relevant personal information to you with a picture. CAs deliver public keys through "certificates," which are generally compliant with the X.509 standard. The CA will publish information to a directory, which contains an entry for each entity with a certificate. It is not too far-fetched to think of the CA as a sort of rigorously maintained cryptographic Yellow Pages.

Public key cryptosystems provide transaction authentication through the use of "digital signatures." Digital signatures are created by the application of a hash function to a piece of data (e.g., a message). This message hash is then encrypted with a sender's private key, as shown in Exhibit 7. The message recipient can use the sender's public key to decrypt the message hash, and rerun the hashing algorithm to make sure the hash has not changed. If the two hashes match, the sender has been properly authenticated. Note that for authentication, the pattern of public/private key use is the reverse of that for confidentiality. For confidentiality, the sender encrypts with the receiver's public key. To provide authenticity, the senders encrypt with their own private key.

Public key cryptography is computationally expensive. For this reason, most modern systems use a combination of public key cryptography and symmetric key cryptography for performance reasons.

Certification services are the support activities needed to verify that the certificates are properly used, to ensure the authenticity and confidentiality of communications and stored data.

The binding of a public key to an entity by a CA does not address all the key management problems associated with asymmetric cryptography. "Key recovery" is another challenge. Data encrypted under a public key

cannot be recovered without the private key. If the private key is rendered inaccessible (through file corruption, token destruction, or failure), it is essential that the cryptosystem owner/operator provide a means for recovering that data.

Another chore associated with key management is *revocation*. In any public key cryptosystem, keys will eventually be compromised, either because they are lost or stolen. Procedures must allow participants to notify an administrator if their keys are compromised, to disseminate the list of compromised keys to all participating entities, and to issue new keys to replace the compromised keys. Because public key binding is typically carried out using x.509 compliant certificates, this process is called *certificate revocation*.

Using public key cryptography requires a "public key infrastructure" or PKI. A PKI provides the administrative structure to manage public key pairs effectively. The PKI functions include key recovery, certificate reissue or renewal, key registration, and key distribution. More complex (and more complete) systems include directory services for registration of participants, distribution of the certificates, a key repository, and a CA hierarchy. A PKI provides the necessary support systems and processes to ensure that all entities are properly bound to their public/private key pairs.

Exhibit 8 illustrates one potential implementation of public key cryptography. The *End User* has a personal private key stored on an encrypted file in the PC or possibly on a more secure device such as a smart card. The client has access to a repository used for *Public Key Storage* to obtain public keys of other entities. The *Web Server* has its own private key stored in a secure file or cryptographic device in the server. The combination of a public and private key pair associated with the *Web Server* and *End User PC* along with cryptographic software on either end enable the two entities to authenticate to each other, send encrypted data, and digitally sign documents. In some cases, an *Authentication Server*, or directory service such as LDAP, can be used to validate the user's current access rights to the system. Authentication services and directories are often used to supplement Certification Revocation Lists to provide faster and more granular authorization. The technology to support directories, certificate revocation, certificate issuance, and other components of the PKI are rapidly advancing and common architectures and methods are still evolving.

A number of specific technologies are available to implement cryptography and certification in today's netcentric environments. If an organization is running a Netscape Web server, it may want to consider Secure Sockets Layer, or SSL, which provides data encryption, server authentication, message integrity, and optional client authentication for a TCP/IP connection. Secure Mime and Pretty Good Privacy, or PGP, are common encryption




Exhibit 8. Implementation of Public Key Cryptography.

solutions for electronic mail. PGP will both authenticate the sender of the message, and encrypt the contents of the message through the use of a public key/private key pair.

In eCommerce solutions, the Secure Electronic Transactions (SET) specification jointly developed by Visa and Mastercard may be considered. SET will require authentication of all parties involved with a credit card transaction through the use of digital signatures and certificates, and will use a separate encryption handshake in order to guarantee both confidentiality and integrity. A number of options are available to implement a PKI, including service providers such as Verisign, and a number of CA products that can be implemented and managed in-house. Elliptic key cryptography can be used to increase the speed of cryptographic functions. The technology choices are rapidly developing, so what was a good solution six months ago may be out of date today. A variety of toolkits that implement these technologies are being rapidly advanced in the marketplace, and the options are evolving as some features are incorporated in to new versions of Web browser, server software, and development environments.

Two useful capabilities provided by cryptographic system include integrity preservation and nonrepudiation:

- *Integrity preservation.* Integrity controls ensure that data maintains a level of quality commensurate with the business needs, and that data is not modified by unauthorized parties or in an unintended fashion. Integrity controls can apply to a single message, an entire transmission, or a piece of data in a database. Digital signatures, message authentication codes, input edits, check digits, checksums, hash values, headers/trailers, and control totals are several methods that are

used to ensure integrity. Access control helps to ensure integrity by allowing access to modify data only by authorized parties. Many controls that assure integrity are outside the scope of this chapter because they would not be classified as security controls. (See the chapter on controls in Section III). The primary controls that are directly related to security are those that involve the implementation of cryptographic functions and those involving access control.

- *Nonrepudiation.* Nonrepudiation means ensuring that one's actions are properly attributed and cannot be denied. This is done with strong audit methods, with entities whose identities have been strongly authenticated. Digital signatures, if properly implemented, can be used to provide nonrepudiation.

Firewalls

Firewalls are used to restrict traffic between networks. They can include a number of features. More conservative firewall technology involves the use of a proxy server that mediates all access through the firewall through a proxy application that is a small, highly trusted piece of code to mediate access through the firewall for a specific service. Less secure, but more flexible approaches use packet filtering to restrict access. These approaches are more flexible because packet filtering rules are extremely flexible; one can write rules to allow or disallow any service, type of packets, etc. Packet filtering is less secure because an organization may allow insecure services through its firewall. Firewalls may also provide authentication based on source location or userID, password, token, or public key authentication. All firewalls have extensive access control, audit logging, and alerting capabilities. Virtual private networks can be created by using encryption-enabled firewalls to establish a secure pipe between several locations over the Internet.

Audit Tools

Compliance auditing is often enforced through a combination of manual and automated procedures. Automated methods are especially important in the netcentric environment because the threats evolve so rapidly that it is only through automation that an organization can truly get a handle on all of the vulnerabilities that may be present across the increasing number of computers involved. These audit tools may scan the network for insecure machines, assess operating system security provisions across a large number of machines, assess the security Web server, or perform other similar functions.

Intrusion Detection

Intrusion detection software is relatively new. Intrusion detection takes audit logging services one step further by monitoring the environment in

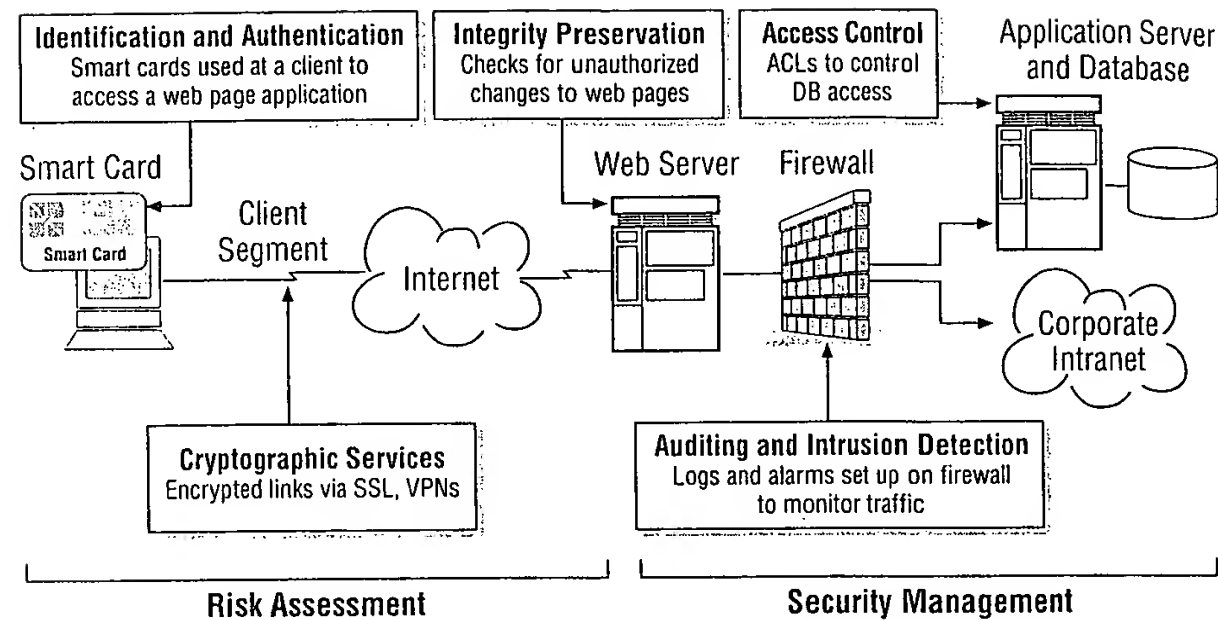


Exhibit 9. Components to Secure a Netcentric Solution.

real time, and actively warning of suspicious activity. Processes continually monitor the system using a set of heuristics, looking for signs of intrusion, such as attempts to take advantage of common security holes, repeated attempts to guess a password, or large concentrated numbers of access list violations. Other intrusion detection devices attach to a network and inspect passing packets for attacks on networking protocols.

Active Security Feature Enhancement Tools

A variety of tools are available to improve security of operating systems and PCs. Some tools, such as HP Virtual Vault and Memco's SEOS, actively improve the security of the UNIX operating system by increasing file protection capabilities, improving sign-on security, and improving segregation of duties. Other tools, such as Finjan, work to combat mobile code threats presented by Java, JavaScript, and Active-X. Still other tools are used to combat viruses and other types of threats. Fraud detection services may be used to catch unauthorized use of credit cards.

Exhibit 9 illustrates a number of components that may be used to secure a netcentric solution.

CONCLUSION

The Internet presents almost limitless opportunity, but comes with a price of almost limitless risk. The security challenges are among the most difficult an organizations will face as it moves into the netcentric environment. The lack of generally accepted methods, the difficulty in integrating consistent

security into the plethora of environments to be protected, and the netcentric characteristics outlined previously will continue to pose significant challenges in the future.

Netcentric technologies are always changing and evolving. On the horizon today we have the next generation of IP (IPng or IPv6), a growing and more complex public key management problem, new types of development environments, and many others. Similar to the development of current Internet technologies, some of these will evolve after a long slow process, and others will burst quickly onto the scene.

Although the enumeration of the security challenges faced in the netcentric world may seem daunting, this technology holds out the promise of better security and enhanced privacy. It is, after all, far easier to forge a written signature on a paper check than it is to forge a digital signature. The eventual universal availability of strong encryption will put even mundane communications beyond the reach of most determined snoops.

Security is only as strong as the weakest link. It is vital that organizations think of security in all its facets, including people, process, and technology components. One key will be to stay abreast of the new technologies and to continually reevaluate the security of the evolving netcentric environment. Security must be considered at every stage of netcentric systems integration. It will be especially important to plan for security upfront, to design security into the architecture and application, to invest in security solutions that will enable the entire enterprise, and to have people with the correct skills developing the security solution.

As an organization selects tools and strategies to maintain its security, it is important to evaluate both leading-edge technologies as well as traditional solutions. The solution that is right for a given circumstance will probably be a mix of the two. To maintain security, it is critical not only to watch out for security bugs in new technologies but for new problems discovered in old technologies, and, in all cases, it is the risk assessment that should drive the efforts.